

(12) UK Patent Application (19) GB (11) 2 365 252 (13) A

(43) Date of A Publication 13.02.2002

(21) Application No 0011201.1

(22) Date of Filing 09.05.2000

(71) Applicant(s)

3Com Corporation
(Incorporated in USA - Delaware)
5400 Bayfront Plaza, M/S 1308, Santa Clara,
California 95052-8145, United States of America

(72) Inventor(s)

Peter I Oliver
Ronald Brown
Mark A Pearce

(74) Agent and/or Address for Service

Brookes Batchellor
102-108 Clerkenwell Road, LONDON, EC1M 5SA,
United Kingdom

(51) INT CL⁷

H04L 12/24 12/26

(52) UK CL (Edition T)

H4K KFMA

(56) Documents Cited

EP 0535270 A1

WO 97/37292 A2

WO 98/25377 A1

WO 93/10495 A1

(58) Field of Search

UK CL (Edition R) H4K KFM, H4P PEUX

INT CL⁷ H04L 12/24 12/26 12/56, H04Q 3/00

Online: EPODOC, JAPIO, WPI

(54) Abstract Title

Network management system user interface which presents graphs of conversation information in response to user selection

(57) A network management system has a user interface 9 which includes areas for listing network devices by name and/or address 17 and displaying a network map 15. A user can select a device or group of devices from the list(s) or map and in response graphs are produced in the display areas 21, 23 and 25. These graphs relate to the communications traffic (ie conversations) between devices. The graphs are typically bar graphs and may be ordered according to the volume of traffic. The graphs may display a summary of the information for a group of devices 21, information for a particular device 23 or information regarding a particular pair of devices 25. The bar graphs may be broken down into colour coded blocks dependent upon the protocol used for the conversation(s).

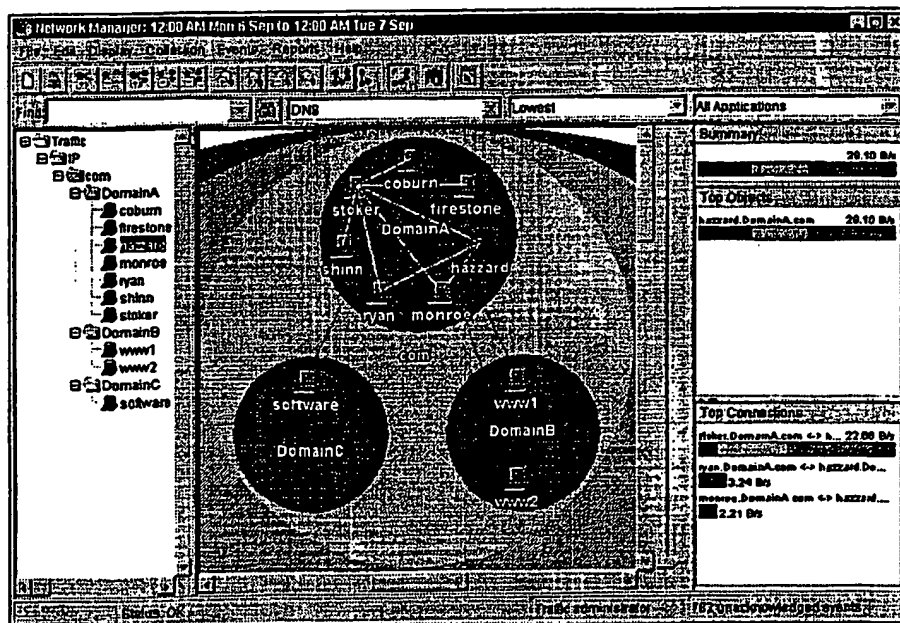


Fig.4

1/10




Fig. 1




Fig.2




Fig.3

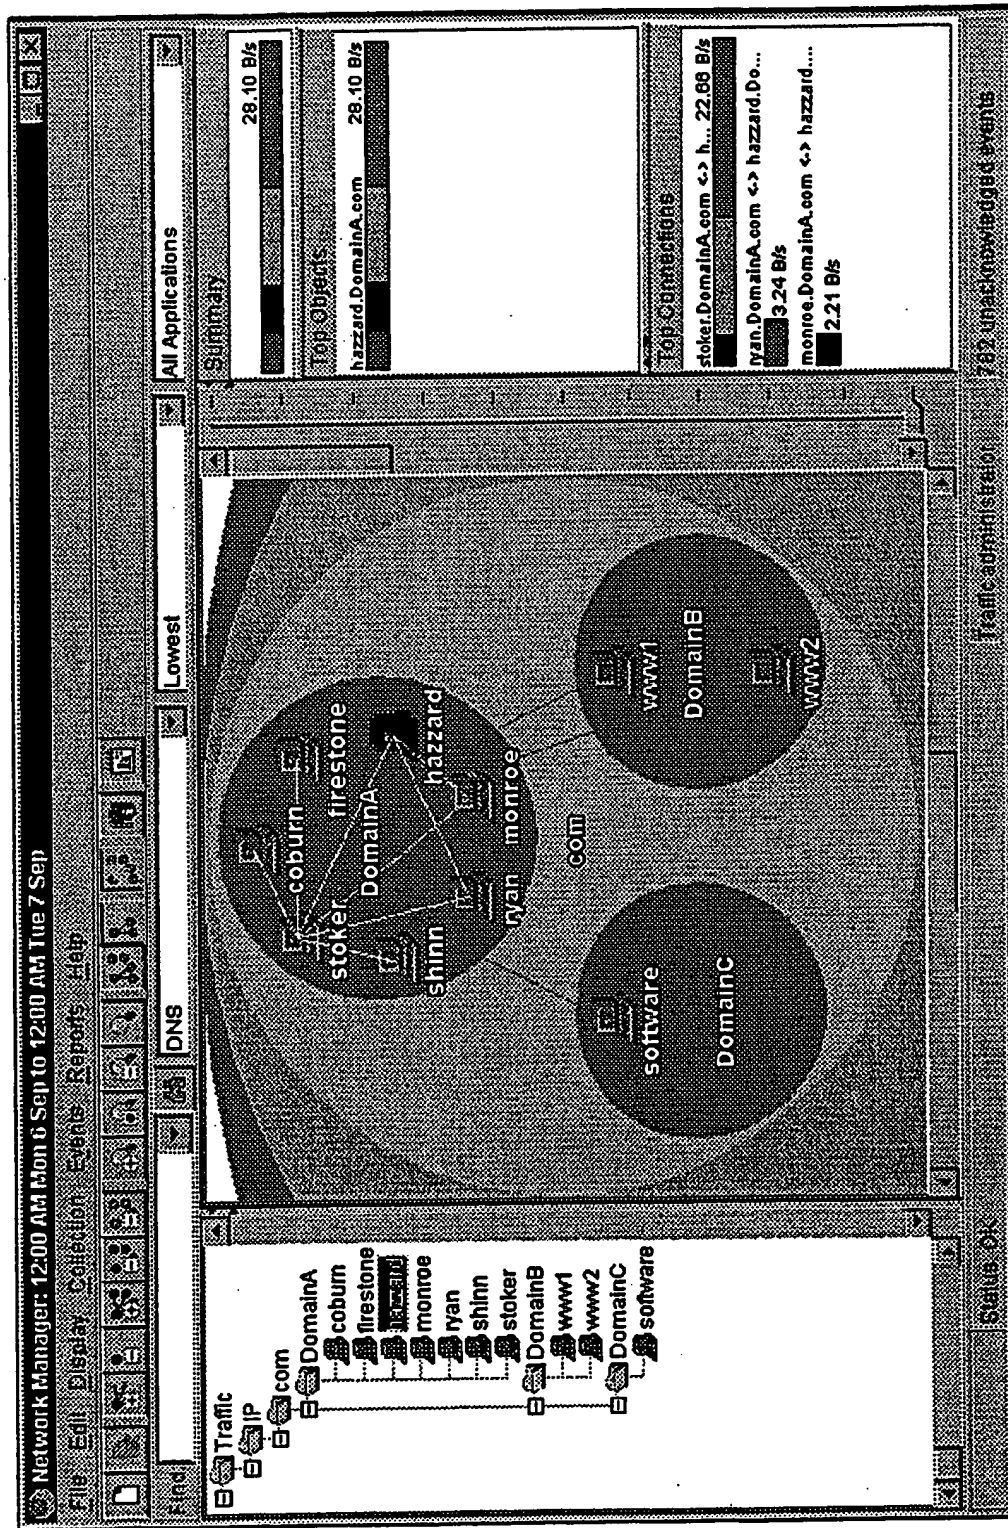




Fig.4





7/10

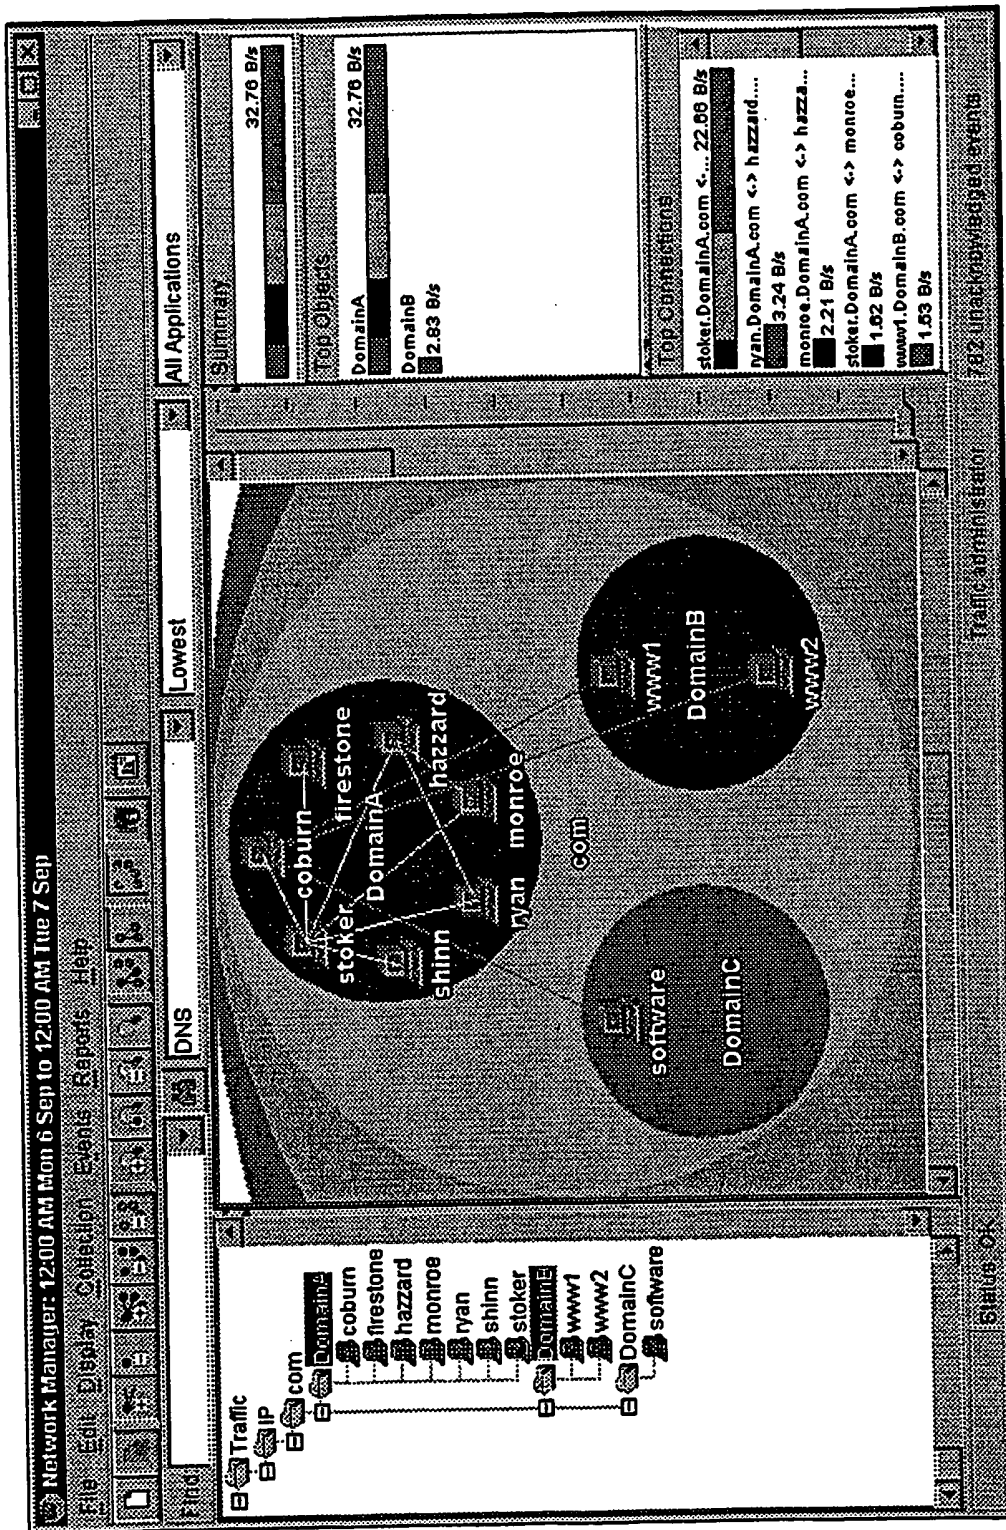




Fig.7





10/10

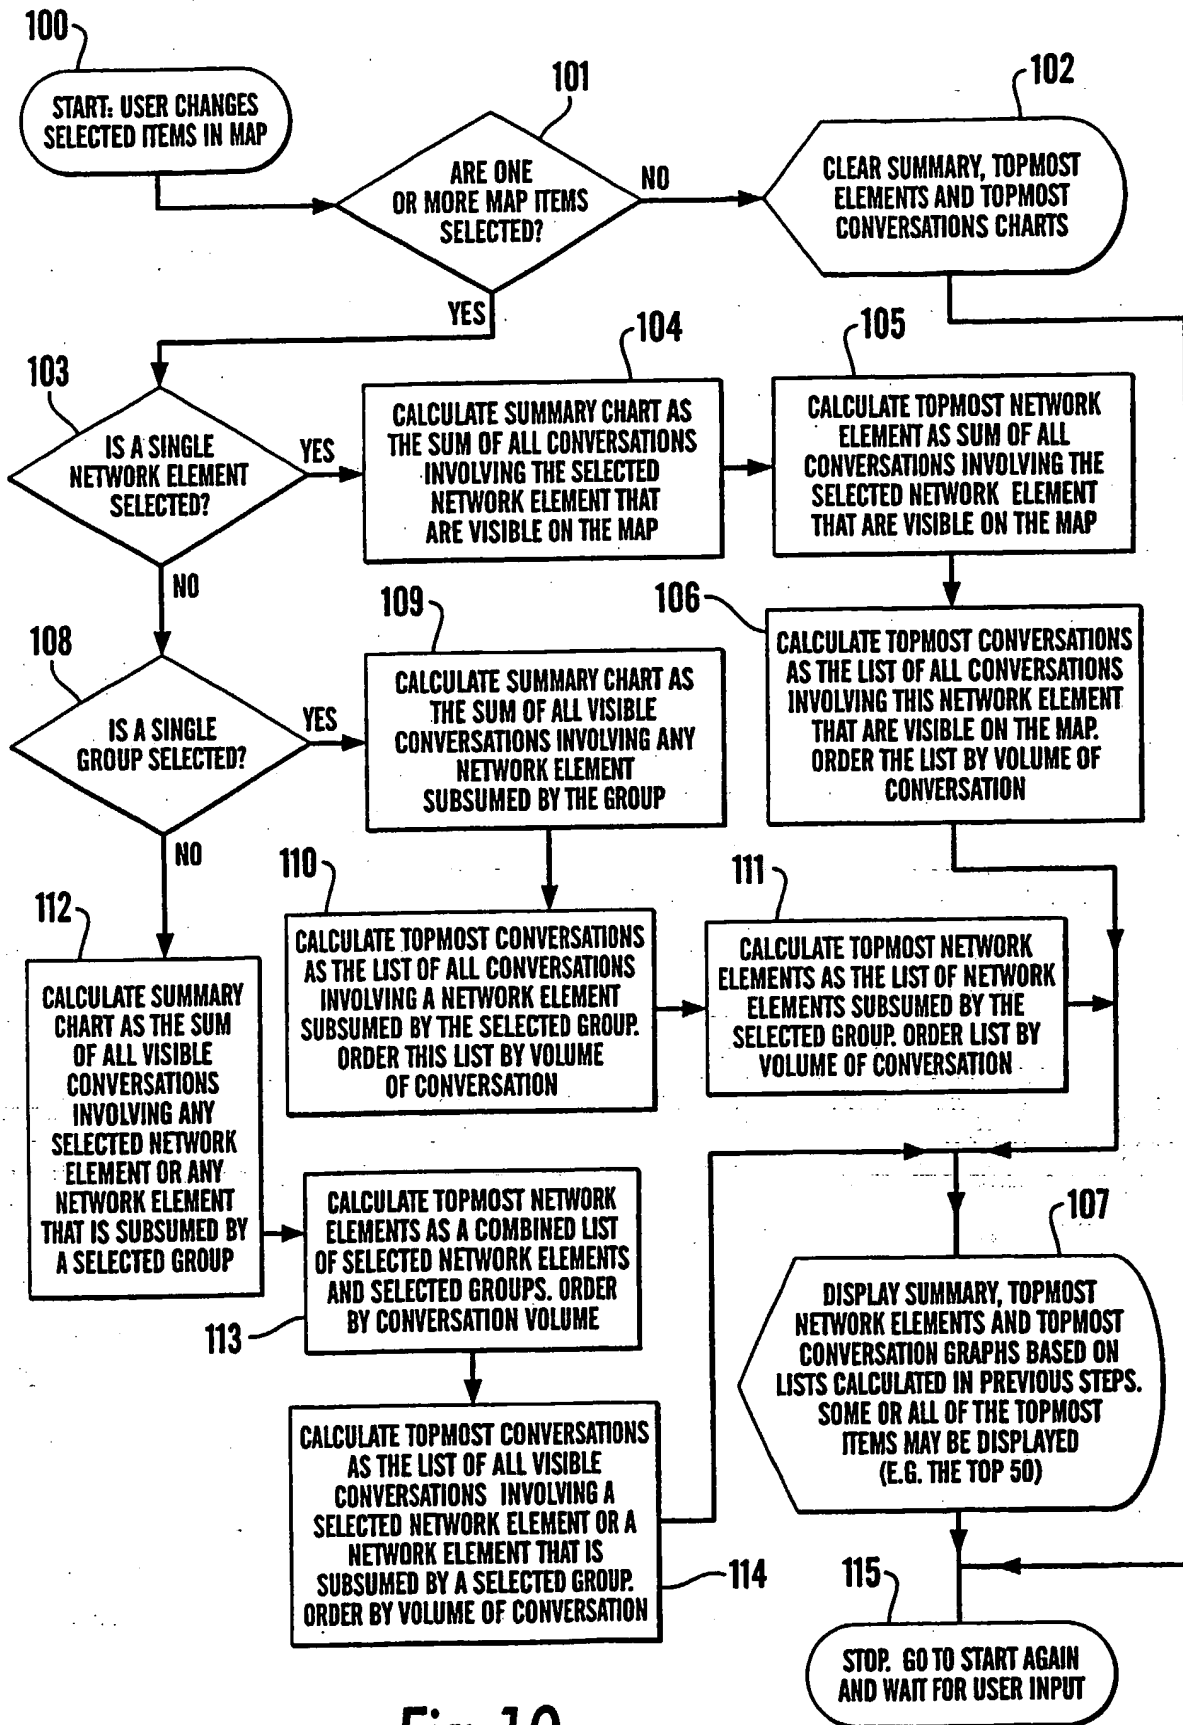


Fig. 10

APPARATUS AND METHOD FOR AUTOMATICALLY PRESENTING
SIGNIFICANT DATA IN RESPONSE TO USER SELECTION IN NETWORK
MANAGEMENT SYSTEMS

5 BACKGROUND OF THE INVENTION

Field of the Invention

 The present invention relates to an apparatus and method for presenting data in response to user selection, and in particular, but not exclusively, to an apparatus and method which graphically presents network management data relating to network devices according to selection by the user.

Description of the Related Art

 Network management systems collect and process data relating to the operation of a network. This application describes management systems for computer networks, such as local area networks (LANs) and wide area networks (WANs), but it will be appreciated that it may be applicable to other types of data communications networks.

 Computer networks typically comprise a plurality of network devices including computers, peripherals and other electronic devices capable of communicating with each other by sending and receiving data packets in accordance with predefined network protocols. Each network device on the network is connected to the network media, which in the case of a LAN network may be coaxial cable, twisted pair cable or fibre optic cable. A network is generally configured with core devices having a plurality of ports, which can be used to interconnect a plurality of media links on the network. Such devices include hubs, switches and routers which pass data packets received at one port to one or more of its other ports, depending upon the type of device.

 A computer network may be managed by a network management system which monitors data traffic, comprising one or more strategically placed monitors or

"probes" on the network. These "probes" monitor data packets transferred across the network and collect data relating to "conversations" between network devices. Whilst the monitoring of conversations is typically carried out by a dedicated monitoring device, the skilled person will appreciate that the monitoring may be carried out by a managed network device such as a managed hub or switch, as is well known in the art.

Such a network management system further comprises a network management station which periodically retrieves the collected data from the "probes", processes the data and stores it in a database or other data storage file. The network management station typically comprises a computer or computer system with the necessary hardware and software to carry out the network management functions, including the necessary memory capacity to store data in the database. The network management station typically presents data from the database to the user by means of a conventional display screen. The data may be presented in a variety of different forms based on selection by the user.

For example, the data may be presented on the display screen in the form of a visual "map" of circles representing groups of related devices with points or icons within the circles representing individual network devices within the groups. Lines between circles and/or icons within the circles represent monitored "conversations". The thickness of the lines may be used to indicate the quantity of data transferred during the conversation and different colours may be used for the lines to indicate the protocol used in the conversation.

In another example, graphs may be used to present information on the display screen about particular conversations selected by the user. For example, data may be presented for conversations involving a selected device or group of devices.

In each case, the user is required to select the manner of presentation of the management information and to identify the devices or groups of devices of interest

and to which the management information should relate.

5 A difficulty with the above described network management system is that the user may spend considerable time deciding which devices or groups of devices are of interest, and may spend time viewing presentations of data which are not significant, or which do not present the data in a useful and understandable manner, before the significant data and most useful manner of presentation of the data is found.

10 The present invention seeks to overcome these shortcomings.

SUMMARY OF THE INVENTION

15 In accordance with a first aspect, the present invention provides a method for presenting data comprising the steps of: generating at least one graph representing data relating to a data object, in response to selection of the data object by a user, the user employing a graphical representation of a plurality of data objects in making the selection, and simultaneously presenting the at least one graph with the graphical representation.

20 In this way, the network user is automatically presented with a graph representing data relating to the selection without having to make further selections for the display of such data.

25 In accordance with a second aspect, the present invention provides a computer program for carrying out the method of the first aspect of the present invention.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a block diagram showing a network management system in accordance with the present invention;

30 Figure 2 is a schematic view of the presentation regions for presentation of management data on a display screen of a network management station in accordance with a preferred embodiment of the present invention;

Figures 3 to 9 are illustrative views of the presentation of management data on a display screen of a network management station in accordance with a preferred embodiment of the present invention, showing different user selections, and

Figure 10 is a flow diagram showing the steps performed by a computer program in accordance with a method according to the present invention.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

Figure 1 shows a network management system 1 in accordance with a preferred embodiment of the present invention. The management system 1 comprises a plurality of monitors or probes 3 situated at strategic locations on a network 5. In the illustrated embodiment, the network 5 comprises three local area networks (LANs), which are managed by the network management system 1. These LANs include devices which are members of domains having the domain names "DomainA.com" indicated by the reference 6, "DomainB.com" indicated by the reference numeral 8 and "DomainC.com" indicated by the reference numeral 10. It will be appreciated that other network configurations are possible. For example, only one probe 3 may be required if it can be positioned at a location to monitor all conversations involving devices on the network, as described below.

The network management system 1 further comprises a network management station 7 having a processor (CPU) 11 for processing management data and memory 12 including random access memory (RAM), a disk drive and a database 14 for storing the management data. The management station 7 further comprises user interfaces 16 such as a display screen 9, a keyboard and a mouse.

The network management system 1 uses the Simple Network Management Protocol (SNMP) to communicate management data. SNMP defines agents, managers and MIBs (where MIB is Management Information Base), as well as various predefined messages and commands for data communication. An agent is present in each managed network device and stores management data, responds to requests from the manager and may send a message to the manager after sensing a

predefined condition. A manager is present within the network management station of a network and interrogates the agents of managed devices on the network using various SNMP commands, to obtain management information. A MIB is a managed object database which stores management data obtained by managed devices, accessible to agents for network management applications. In the preferred embodiment, the MIBs employed are RMON (Remote Monitoring MIB as defined in RFC1271) and RMON2 (Remote Monitoring MIB version2 as defined in RFC2021)

It will be appreciated that the present invention is not restricted to management using SNMP and data sources other than those described below may be used.

In accordance with the preferred embodiment, probes 3 include SNMP agents which collect RMON and/or RMON2 data. Management station 7 includes an SNMP manager which periodically retrieves the RMON/RMON2 data from all the probes 3, for example, every 30 minutes, for conversations which have occurred during a preceding time interval, as explained below. The management station 7 processes the data and stores it in the database 14, using conventional techniques, for example as described in UK Patent Application GB 2 337 903A, which is incorporated herein by reference.

More specifically, the management station 7 stores data about "conversations" (i.e. data communications) between two devices on the network which have been "seen" by the probes 3, hereinafter called "conversation data". Conversation data includes, for each conversation, the identity of the two devices involved in the conversation, the start time and end time of the conversation, the quantity of data transferred during the conversation and the protocols used. This information is obtained by probes 3 through observing the data packets transferred during the conversation. For example, the header of each data packet will contain an address for the source device and an address for the destination device. The protocols employed can similarly be obtained from information in the data packets. The number of data

packets transferred during each conversation can be counted by the probe. The start time and end time of the conversation can be noted by the probe and, in a preferred embodiment a "timestamp" applied to the conversation data to indicate the time interval (for example the hour or half hour) in which the conversations occurred. The timestamp is then used for identifying the data for conversations in the relevant time interval when the management station 7 retrieves the data from the probes 3. Conversation data is indexed in the database by *inter alia* the device identifier ("Device ID" or "Object ID") for the two devices involved in the conversation, which may be an identifier for an individual device involved in the conversation or the identifier for a group of devices which includes an individual device involved in the conversation.

The network management station 7 may be operated to display on its display screen 9 selected conversation data from the database 14. For example, the user may use the keyboard or mouse to select data for monitored conversations over a specified time interval, which may be the time period for which the management station 7 has been monitoring the network 5 or may be based on a daily, weekly or monthly time interval. Multiple time intervals may be selected by the user. In accordance with the preferred embodiment, the data retrieved from the probes 3 for each time interval is stored in a separate record in the database. This means that when the user makes a selection of one or more time intervals, the management station can copy the records falling within each specified time interval from the database 14 and store the records in a first cache area of the memory 12 for use during management operations. In particular the cached data is used in calculations to generate an appropriate display of the data as described below.

In accordance with the present invention, network management data is simultaneously displayed in a plurality of different graphical forms.

Figure 2 schematically illustrates the display area of display screen 9 of the network management station 7 in accordance with the preferred embodiment. The

display area includes a toolbar area 19, a map area 15, a list area 17 and a graph area 20.

The main, centrally located display area, map area 15, is used to display a "map" representing the traffic on the network 5, in which circles represent groups of related network devices, icons within the circles represent individual network devices within the groups and lines between circles and/or icons represent monitored conversations.

The list area 17 is used to display a list of names of network devices shown on the map in map area 15. It will be appreciated that this area may list the devices by other identifiers such as IP address. The list area 17 is shown to the left hand side of the map in the display area of Figure 2.

The graph area 20 comprises three graph regions, and is located on the right hand side of the map in the display area of Figure 2. The graph area 20 is used to display graphs of data relating to network devices selected by the user using the map in the map area 15 or the list in the list area 17. The graph regions in the graph area 20 are Summary Region 21, Top Objects Region 23, and Top Connections Region 25. In accordance with the preferred embodiment, the network management station 7 automatically generates and presents graphs in each of these regions according to the device or group of devices selected by the user from the list region 17 and/or the map region 15. Thus the user is automatically presented with useful graphs relating to the selected devices or groups of devices.

The graphs generated in each of the graph regions will now be described, and the manner of generation of the graphs, using the data in the database 14, in accordance with the preferred embodiment of the present invention will then be described.

In accordance with the preferred embodiment the graphs comprise bar charts which present the management data and more particularly the conversation data for a selected time period. It will be appreciated that other forms of graph may be used to

present the data, and the present invention is not limited to the use of bar charts. The bar charts comprise bars, each bar representing a particular network device or group of network devices and the length of each bar represents the quantity of data packets sent and received by the particular device or group during monitored conversations over a prespecified time period. More specifically, each bar is a stacked bar comprising a series segments stacked together in a linear array to form the bar. Each segment of a bar represents the quantity of data packets sent and received by the particular device or group represented by the bar using a particular network protocol. The segments within a bar may be distinguished from each other by the use of different colours for different protocols. The sum of the segments in a particular bar thus represents the total quantity of data packets sent and received by the particular device or group of devices during conversations in the specified time period.

Summary Region

The bar chart presented in the Summary Region 21 represents data for all the conversations in the specified time period for the device or group of devices selected using the map or list. As illustrated, for example in Figure 4 in which a single device named "hazzard.DomainA.com" is selected, the graph is a bar chart comprising a single stacked bar that displays the sum of all of the data traffic entering and leaving the selected device "hazzard.DomainA.com". Each element or segment of the stacked bar represents an individual network protocol or named group of network protocols used in monitored conversations involving the device "hazzard.DomainA.com", and the length of each segment represents the quantity of data (in bytes, as indicated above the bar) transferred in conversations using the relevant protocol i.e. the volume of traffic involving the devices.

Top Objects Region

The bar chart presented in the Top Objects Region 23 represents the conversations involving each device within the selected device or group of devices. Each bar represents the monitored conversations involving a particular device within the device or groups of devices selected using the map or list. The bars are presented

in the bar chart in order by volume (i.e. length of bar) with the longest bar at the top of the bar chart and the shortest bar at the bottom. Thus, the most active, and therefore significant, devices in terms of conversation volume are presented at the top of the bar chart. The number of bars in the bar chart in the Top Objects Region is limited, in the preferred embodiment to 50, so that data for devices within the selection which are comparatively inactive, and therefore less significant to the user, are omitted from the chart. It will be appreciated that the number of bars displayed may be chosen by the user and is typically in the region of 10 to 100. This graph is a stacked bar chart that contains multiple bars; each bar represents a "significant device" or collapsed group (as discussed in more detail below). As with the bar chart in the Summary Region 21, each element or segment of each stacked bar represents an individual network protocol or named group of network protocols.

The contents of the bar chart in the Top Objects Region 23 change as follows depending on the context of the selection.

When a single device is selected, as shown in Figure 4 (or a collapsed group is selected as shown in Figure 9) the bar chart simply contains the sum of all of the traffic entering, leaving, or contained within the selected device (or collapsed group). Thus in Figure 4, where the device "hazzard.DomainA.com" is selected, the bar chart contains a single bar representing all the conversations involving "hazzard". In other words it is a duplicate of the bar chart in the Summary Region.

When a single group of devices (not collapsed) is selected, as shown in the Figure 5, the bar chart shows the conversations involving the most active devices contained within the selected group, ordered by volume of conversation data. Each bar represents the sum of the traffic entering or leaving each "child" device within the selected group. Thus, in Figure 5, where the group of devices "DomainA" are selected, the bars represent the individual devices within the group, of which the top device, which has involved the transfer of the most conversation data is "hazzard", and the next most active device in terms of conversation data volume is "stoker".

Thus, it can be seen that this graph allows the user to rapidly identify the devices within the group which are consuming the most network resources.

If multiple devices are selected, as shown in Figure 6 (or multiple groups are selected as shown in Figure 7), each bar in the bar chart represents the sum of data traffic entering, leaving, or contained within a single device (or group) from the selection. The bars are sorted by volume of traffic. Thus, as shown in Figure 6, in which the devices "hazzard", "stoker" and "ryan" are selected, each device is represented by a single bar, thus the bar chart includes three bars revealing information relating to the data traffic involving each of the devices. As will be appreciated, selecting multiple devices and/or groups allows the user to quickly compare selected devices and/or groups. Thus, it can be seen from Figure 6 that the device "ryan" is consuming considerably fewer network resources compared to the active devices "hazzard" and "stoker".

Finally, if a combination of devices and groups is selected, as shown in Figure 8, each bar in the bar chart represents the sum of data traffic entering, leaving, or contained within a selected device or a selected group. The bars are sorted by volume of traffic. Thus, as shown in Figure 8, in which the device "www.DomainB.com" is selected in combination with the group "DomainA", the device "www" is represented by a single bar, and the group "DomainA" is represented by a single bar. Thus the bar chart includes only two bars revealing information relating to the data traffic involving the selected device and the selected group.

Top Connections Regions

The graph presented in the Top Connections Region 25 represents the connections between each selected network device, or each network device within a selected group of network devices, and other network devices on the network, which are carrying the most network traffic. Thus, each bar represents the monitored conversations on a particular "connection" i.e. conversations between two network devices, at least one of which is within the selection. This graph comprises stacked bars in a bar chart, each bar representing a "significant" visible network connection,

and each element of each stacked bar representing the individual network protocol(s). In accordance with a preferred embodiment of the present invention, a "significant" connection is defined as one which consumes a high volume of network traffic and the bars within the bar chart are sorted by order of volume of network traffic (highest volume at top). As with the graph in the Top Objects Region, in the preferred embodiment, only the 50 busiest connections that are attached to selected devices and/or groups are presented in the bar chart.

The contents of this bar chart in the Top Connections Region 25 change as follows depending on the context of the selection.

When a single device is selected, as shown in Figure 4 (or a collapsed group is selected as shown in Figure 9, and described below) the bar chart includes a bar for each of the devices which have been connected to the selected device during monitoring. Thus in Figure 4, where the device "hazzard" is selected, the bar chart contains bars representing conversations with "hazzard", "stoker" and "monroe". These three bars in the bar chart represent all the conversations involving "hazzard". It will be appreciated that for devices having conversations with a large number of devices, only the 50 busiest connections, in terms of volume of data traffic, would be presented.

When a single group of devices (not collapsed) is selected, as shown in the Figure 5, the bar chart shows the conversations involving the devices within the selected group and other devices on the network, ordered by volume conversation data. Each bar represents the traffic transferred between each "child" device within the selected group and another device. Thus, in Figure 5, where the group of devices "DomainA" is selected, the bars represent connections involving devices within the group "DomainA". The top bar illustrates the connection between "hazzard" and "stoker", showing that conversations between these two devices have involved the transfer of the most conversation data, and the bottom bar illustrated in the view of Figure 5 shows the connection between "software.DomainC.com" and

"coburn.DomainA.com" which has involved the transfer of considerably less conversation data.

5 If multiple devices are selected, as shown in Figure 6 (or multiple groups are selected as shown in Figure 7), each bar in the bar chart represents the data traffic transferred between a "child" device within the selected group and another device, with the bars sorted by volume of traffic. Thus, as shown in Figure 6, in which the devices "hazzard", "stoker" and "ryan" are selected, data relating to conversations involving each device and another device is represented by a single bar. Thus, as with
10 Figure 5, the top bar illustrates the connection between "hazzard" and "stoker", since monitored conversations between these two devices have involved the transfer of the most conversation data. The next bar represents conversations involving the devices "ryan" and "hazzard" which have involved considerably less data traffic.

15 Finally, if a combination of devices and groups is selected, as shown in Figure 8, each bar in the bar chart represents the data traffic transferred between a selected device or a "child" device within a selected group and another device, with the bars sorted by volume of traffic. Thus, as shown in Figure 8, in which the device "www.DomainB.com" is selected in combination with the group "DomainA", data
20 relating to conversations involving "DomainB" and each device within the group "DomainA" and another device is represented by a single bar. Thus, as with Figure 5, the top bar illustrates the connection between "hazzard" and "stoker", which are both "child" devices of the selected group "DomainA" since monitored conversations between these two devices have involved the transfer of the most conversation data.
25 The next bar represents conversations involving the devices "ryan" and "hazzard", also within the group "DomainA" which have involved considerably less data traffic.

Collapsed Group

30 Figure 9 illustrates the selection is a "collapsed" group of devices. A collapsed group is one in which all of the data relating to "child" devices within the group has been aggregated and represented as a single entity on the map. When data

relating to groups of devices are collapsed, in response to user request, the management station 7 performs a calculation to aggregate the data relating to the group of devices for presentation on the map in map area 15 and in the associated graphs in graph area 20. Specifically, the data for the selected time intervals, which is
5 stored in a first cache area of memory 12, as previously described, is aggregated to relate to the group and stored in a second cache area in memory. This data in the second cache area thus relates to "group to group" conversations, in contrast to data in the first cache area which relates to "device to device" conversations. As will be appreciated, the data in the first cache area, and in database 14, remains unaffected by
10 the collapsing of devices, thus enabling the user to reverse the collapsing so that devices within the collapsed group can be represented individually on the map, and their corresponding data displayed in the graph area 20.

As explained above, collapsed groups are treated as single devices for the
15 generation of the graphs in the Summary Region 21, the Top Objects Region 23 and the Top Connections Region 25. Thus, in Figure 9, in which the collapsed group "DomainA" is selected, the Summary Region 21 represents data for all the conversations involving "child" devices within the selected collapsed group of devices "DomainA". The graph is a bar chart comprising a single stacked bar that displays the
20 sum of all of the data traffic entering, leaving and contained within the selected collapsed group "DomainA". Each element or segment of the stacked bar represents an individual network protocol or named group of network protocols used in monitored conversations involving the collapsed group "DomainA", and the length of each segment represents the quantity of data (in bytes, as indicated above the bar)
25 transferred in conversations using the relevant protocol i.e. the volume of traffic involving the devices. The graph in the Top Objects Region 23 is the same, since the collapsed group "DomainA" is treated as a single device, and thus the "child" devices within the group are not treated individually. Finally, the graph in the Top Connections Region 25 represents conversations between the collapsed group
30 "DomainA" and external devices. Since the only external devices represented in the map are also collapsed groups, namely the groups "DomainB" and "DomainC", the

two bars shown in the graph in the Top Connections Region of Figure 9 related to conversations between the collapsed groups "DomainA" and "DomainB" and "DomainA" and "DomainC".

5 Method of the Preferred Embodiment

Figure 10 is a flow diagram illustrating the method used by the network management station 7 for presenting management data and in particular conversatin data as illustrated in Figures 2 to 9 in accordance with a preferred embodiment of the present invention. The method steps are preferably carried out by a computer
10 program running in the network management station 7, but it will be appreciated that in other embodiments the method may be implemented by other means such as in hardware.

The program of Figure 10 generates graphs in the graph area 20 in response to
15 selection of a device or group of devices by the user from the map or list. Such selection is performed by the user using the user interfaces 16, for example by a predetermined keystroke using a keyboard or by clicking on a mouse over the representation of the device on the map/list in the display area of display screen 9.

20 The program starts at step 100 in response to change in the selection by the user. At step 101, the program considers whether there are one or more devices or groups of devices selected on the map. If step 101 determines that there is no selection the program continues with step 103, in which the program clears the graph area 20 of the display area of display screen 9 by presenting a blank area in place of
25 any existing graphs (as shown in Figure 3), and the program ends at step 115. Alternatively, if step 101 determines that a selection exists, the program continues with step 103.

30 At step 103, the program considers whether a single element (i.e. a single device or collapsed group of devices) has been selected. If step 103 determines that a

single element is selected, the program continues with program steps 104 to 106, as described below. Otherwise the program continues with step 108.

5 At step 108, the program considers whether a single group of devices (not collapsed) has been selected. If step 108 determines that a single group of devices is selected, the program continues with program steps 109 to 111, as described below. Otherwise the program determines that multiple devices/groups are selected, and the program continues with steps 112 to 114, as described below.

10 Steps 104 to 106 are performed when a single element (i.e. a single device or a single collapsed group) is selected, as shown by way of example in Figure 4 and Figure 9. In step 104, the program retrieves data for all the monitored conversations involving the selected element. If a time interval has been selected by the user, this data is retrieved from the first or second cache area of the memory as described
15 above. Using this data, the program calculates the sum of the volume of data traffic for different network protocols for the conversations and generates a stacked bar chart, as previously described, representing the total data traffic for presentation in the Summary Region 21.

20 In step 105, the program reproduces the stacked bar chart generated in step 104 for the Top Objects Region 23.

25 In step 106, the program sorts the conversation data involving the selected single device according to the other device involved in the conversation. The sorted data thus relates to a connection between the selected device (or collapsed group) and another single device (or collapsed group). The sorted data is then ordered according to traffic volume for each connection of the 50 connections involving the most data traffic, and the program generates a bar, ordered according to data traffic volume, for presentation in a bar chart in the Top Connections Region 25. Once steps 104 to 106
30 are completed, and three graphs have been generated, the program continues with step 107.

In step 107, the program presents the generated graphs in the appropriate regions on the display screen 9. In the case of the graphs in the Top Objects and Top Connections Regions, where there may be insufficient room in the "window" associated with the graph region to display all the bars in the bar charts, the top part of the chart is presented, and the bar chart can be scrolled using a scroll bar associated with the window of the graph region 20.

Steps 109 to 111 are performed when a single group of elements is selected, as shown by way of example in Figure 5, in which a single group of devices is selected. It will be appreciated that the single group may involve collapsed groups, for example if the single group "com" were selected from the map of Figure 9.

In step 109, the program retrieves data for all the monitored conversations involving devices within the selected group of elements. If a time interval has been selected by the user, this data is retrieved from the cache area of the memory 12 as described above. Using this data, the program calculates the sum of the volume of data traffic for different network protocols for the conversations. The program uses this calculation to generate a stacked bar chart as previously described, the bar chart representing the total data traffic involving the group (i.e. between elements in the group and between an element in the group and an element outside the group) for presentation in the Summary Region 21.

In step 110, the program sorts the conversation data involving the selected group of elements (devices and/or collapsed groups) by element. Thus for a selected group of devices, the data is sorted according to the "child" device within the group involved in the conversation. For the selection of a group including collapsed groups, the data is sorted by collapsed group and not child device. The sorted data for each element therefore relates to data traffic for a particular "child" device within a group or data traffic for a collapsed group. The sorted data is then ordered according to traffic volume and, for each of the 50 elements (devices or collapsed groups)

involving the most data traffic, the program generates a bar, for presentation in a bar chart in the Top Objects Region 23 and ordered according to data traffic volume.

5 In step 111, the program sorts the conversation data involving the selected group of elements (devices and/or collapsed groups) according to the connection, i.e. according to the two elements involved in the conversation. The sorted data for each connection relates to data traffic across a connection between an element in the selected group and another single element (single device or collapsed group). The sorted data for the connection is then ordered according to traffic volume and for each
10 of the 50 connections involving the most data traffic, the program generates a bar, for presentation in a bar chart in the Top Connections Region 25 and ordered according to data traffic volume. Once steps 109 to 111 are completed, and three graphs have been generated, the program continues with step 107, as described above.

15 Steps 112 to 114 are performed, by default, when neither a single element (device or collapsed group) nor a single group is selected i.e. when multiple devices and/or collapsed groups and/or groups of devices are selected. Examples of this type of selection are shown in Figures 6, 7 and 8.

20 For this type of selection, a group of devices is treated as a collapsed group (i.e. the conversation data is aggregated, as previously described) in the preferred embodiment. However, it will be appreciated that in other embodiments it could be treated in the same way as in the previously described selections.

25 In step 112, the program retrieves data for all the monitored conversations involving the selected multiple devices, collapsed groups and groups. If a time interval has been selected by the user, this data is retrieved from the first or second cache area of memory as described above. Using this data, the program calculates the sum of the volume of data traffic for different network protocols for the conversations
30 and generates a stacked bar chart, as previously described, representing the total data traffic for the selected multiple devices, collapsed groups and groups of devices for

presentation in the Summary Region 21.

In step 113, the program sorts the conversation data involving the selected multiple devices, collapsed groups and groups of devices according to element i.e. single devices and collapsed groups involved in the conversation. The sorted data for each element therefore relates to data traffic involving a particular selected single device or collapsed group (which could be a selected collapsed group or a selected group of devices). The sorted data is then ordered according to traffic volume and for each of the 50 elements involving the most data traffic, the program generates a bar, ordered according to data traffic volume, for presentation in a bar chart in the Top Objects Region 23.

In step 114 the program sorts the conversation data involving the selected multiple devices, collapsed groups and groups of devices according to connection, i.e. according to the two elements involved in the conversation. The sorted data for each connection therefore relates to data traffic across a connection between one of the selected devices or collapsed groups and another single element (single device or collapsed group). The sorted data is then ordered according to traffic volume and for each of the 50 connections involving the most data traffic, the program generates a bar, ordered according to data traffic volume, for presentation in a bar chart in the Top Connections Region 25. Once steps 112 to 114 are completed, and three graphs have been generated, the program continues with step 107, as described above.

The computer program of the preferred embodiment may be provided on a computer readable medium such as a magnetic or optical disk which may be loaded into the disk drive of the network management station. Alternatively, the computer program may be carried on a computer system having a website which permits downloading of the computer program over the Internet on a carrier wave to the network management station.

It will be appreciated that various modifications and changes may be made to the described embodiment. For example, whilst the graphs of the described embodiments indicate data quantity for monitored conversations they could equally well represent the time duration of monitored conversations. In addition the data could be sorted in manners other than by volume order, such as according to protocol or time. Also, the data could be stored in the database or other data storage according to the time of collection from the probes rather than according to a timestamp indicating the time interval in which the conversation occurred. It is intended to include all such variations, modifications and equivalents which fall within the spirit and scope of the present invention as defined in the accompanying claims.

CLAIMS:

1. A method for presenting data comprising the steps of: receiving an indication of the selection of a data object by a user, the user employing a graphical representation of a plurality of data objects having associated data in making the selection, and automatically generating at least one graph representing data relating to the selected data object in response to the received indication.
5
2. A method as claimed in claim 1, wherein prior to the step of generating, the method further comprises the step of: retrieving data from a data storage, wherein the data is stored in the data storage with a corresponding time.
10
3. A method as claimed in claim 2, wherein the at least one graph comprises a bar chart showing all available data relating to the selected data object for a prespecified time period.
15
4. A method as claimed in claim 3, wherein the data comprises network management data and each of the plurality of data objects relate to an individual network device or a group of network devices.
20
5. A method as claimed in claim 1, wherein the data comprises network management data and each of the plurality of data objects relate to an individual network device or a group of network devices.
6. A method as claimed in claim 5, wherein the at least one graph comprises a bar chart showing the data associated with each network device within the selected data object.
25
7. A method as claimed in claim 6, wherein the bar chart comprises a bar representing data for each network device.
30

8. A method as claimed in claim 5, wherein the network management data relates to monitored conversations involving the plurality of data objects, the data for each conversation including the volume of data transferred during the conversation and the identity of two devices involved in the conversation.

5

9. A method as claimed in claim 8, wherein the at least one graph comprises a bar chart, each bar showing the data relating to conversations between two network devices, wherein the selected data object includes at least one of the two network devices.

10

10 A method as claimed in claim 1, wherein prior to the step of receiving an indication, the method further comprises the step of: retrieving data for a prespecified time period, and caching the data in a cache storage.

15

11. A method as claimed in claim 10, wherein after the step of receiving the indication, the method further comprises the step of: extracting data from the cache storage relating to the selected data object, and using the extracted data to generated the at least one graph.

20

12. A method as claimed in claim 11, wherein the data comprises network management data and each of the plurality of data objects relate to an individual network device or a group of network devices; the network management data relating to monitored conversations involving the plurality of data objects, the data for each conversation including the volume of data transferred during the conversation and the identity of two devices involved in the conversation;

25

wherein the at least one graph comprises:

a summary graph comprising bar chart showing the data associated with each network device within the selected data object;

a device graph comprising a bar chart showing the data associated with each network device within the selected data object, and

30

a conversation graph comprising a bar chart, each bar showing the data relating to conversations between two network devices, wherein the selected data object includes at least one of the two network devices.

5 13. A method as claimed in claim 12, further comprising the step of: simultaneously presenting the at least one graph with the graphical representation.

10 14. A method as claimed in claim 13, wherein the device graph comprises a bar for each of a predetermined number of network devices, the method further comprising the steps of: for each network device, determining the sum of the volume of data transferred during conversations involving the network device, and selecting for inclusion in the device graph the predetermined number of network devices having the largest sums.

15 15. A method as claimed in claim 13, wherein the conversation graph comprises a bar for each of a predetermined number of network connections between two network devices, the method further comprising the steps of: for each network connection, determining the sum of the volume of data transferred during conversations involving the two network devices, and selecting for inclusion in the conversations graph the
20 predetermined number of network connections having the largest sums

16. A method as claimed in claim 1, further comprising the step of: simultaneously presenting the at least one graph with the graphical representation.

25 17. A method as claimed in claim 1, wherein the graphical representation used to select the data object is selected from the group consisting of: a list of the names of the plurality of data objects; a list of the addresses of the plurality of data objects, and a map illustrating the plurality of data objects.

30 18. A method as claimed in claim 1, wherein the graphical representation is presented on a display screen and the indication of the selection is provided in

response to user interaction selected from the group consisting of: a keystroke; a mouse click or pressure applied to the screen.

19. A computer program for presenting data in a user operable system in which a user employs a graphical representation of a plurality of data objects for selecting a data object, the plurality of data objects having associated data for presentation, the program comprising: a program step for receiving an indication of the selection of a data object by a user, and a program step for automatically generating at least one graph representing the associated data for the selected data object in response to the received indication.



Application No: GB 0011201.1
Claims searched: 1-19

Examiner: Owen Wheeler
Date of search: 23 November 2000

Patents Act 1977
Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:
UK CI (Ed.R): H4K (KFM) H4P (PEUX)
Int CI (Ed.7): H04L: 12/24, 12/26, 12/56; H04Q: 3/00
Other: Online: EPODOC, JAPIO, WPI

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
X	EP 0535270 A1 [DIGITAL SYSTEMS INTERNATIONAL] See Figs. 4,8 and 9 and page 9 line 30- page 11 line 7.	1-3,18,19
A	WO 98/25377 A1 [NORTHERN TELECOM]	
A	WO 97/37292 A2 [CABLETRON]	
X	WO 93/10495 A1 [CABLETRON] See Fig. 9 and page 36 final paragraph to page 38 first paragraph.	1-7,10, 11,16-19

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.

THIS PAGE BLANK (USPTO)

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

☐ **BLACK BORDERS**

☐ **IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**

☒ **FADED TEXT OR DRAWING**

☒ **BLURRED OR ILLEGIBLE TEXT OR DRAWING**

☐ **SKEWED/SLANTED IMAGES**

☐ **COLOR OR BLACK AND WHITE PHOTOGRAPHS**

☐ **GRAY SCALE DOCUMENTS**

☐ **LINES OR MARKS ON ORIGINAL DOCUMENT**

☐ **REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**

☐ **OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.

THIS PAGE BLANK (USPTO)